An Algorithm for L_∞ Approximation by Step Functions

Quentin F. Stout
Computer Science and Engineering
University of Michigan
Ann Arbor, MI 48109-2121
qstout@umich.edu +1 734.763.1518

arXiv:1412.2379v2 [cs.DS] 4 May 2015

Abstract

An algorithm is given for determining an optimal b-step approximation of weighted data, where the error is measured with respect to the L_∞ norm. For data presorted by the independent variable the algorithm takes $\Theta(n + \log n \cdot b(1 + \log n/b))$ time and $\Theta(n)$ space. This is $\Theta(n \log n)$ in the worst case and $\Theta(n)$ when $b = O(n/\log n \log\log n)$. A minor change determines an optimal reduced isotonic regression in the same time and space bounds, and the algorithm also solves the k-center problem for 1-dimensional weighted data.
Keywords: step function approximation; reduced isotonic regression; variable width histogram; weighted k-center; interval tree of bounded envelopes

1 Introduction

Step functions are a fundamental form of approximation, arising in variable width histograms, databases, segmentation, approximating sets of planar points, piecewise constant approximations, etc. Here we are interested in L_∞ stepwise approximation of weighted data. By weighted data $(\mathbf{y}, \mathbf{w})$ on $1 \ldots n$ we mean values (y_i, w_i), $1 \le i \le n$, where y_i is an arbitrary real number and w_i (the weight) is a positive real number. For integers $i \le j$ let $[i : j]$ denote $i \ldots j$. A function f on $[1 : n]$ is a *b-step function* iff there are indices $j_1 = 1 < j_2 < \ldots < j_{b+1} = n + 1$ and real values C_k, $k \in [1 : b]$, such that $f_i = C_k$ for $i \in [j_k : j_{k+1} - 1]$. f is an *optimal L_∞ b-step approximation of* $(\mathbf{y}, \mathbf{w})$ iff it minimizes the weighted L_∞ error, $\max\{w_i \cdot |f_i - y_i| : i \in [1 : n]\}$, among all b-step functions. Since a step can be split into smaller ones, we do not differentiate between "b steps" and "no more than b steps".

Many algorithms have been developed for L_∞ b-step regression [2, 3, 4, 5, 6, 8, 9, 11, 12, 13, 14]. The first $\Theta(n \log n)$ time deterministic algorithms [6, 11] were decidedly impractical, relying on parametric search. A more feasible $\Theta(n \log n)$ algorithm appeared in [3]. However, the time of these algorithms does not improve when b is small, which is the typical case of interest. We present a faster algorithm that is $\Theta(n + \log n \cdot b(1 + \log n/b))$ when the data is presorted by independent coordinate.

2 L_∞ b-Step Approximation

At a high-level overview, our algorithm shares aspects of those in [2, 5, 9], with important differences:

1. Build an interval tree to determine the regression error of an arbitrary interval if it is a single step.
2. Use "search in a sorted matrix" to find the minimal possible error for a b-step approximation.

The search uses a feasibility test which is given ϵ and decides if there is a b-step approximation with error $\le \epsilon$. If there is such an approximation then the test produces one. We incorporate important improvements to this approach: feasibility tests are used during tree construction, not just during the search; tests do not determine the minimal regression error of an interval, merely that it is sufficiently small or too large; previous searches, except the randomized version in [13], were not search in a sorted matrix; and we exploit the fact that calculations at one stage of the search are related to those of the previous stage. We will show:

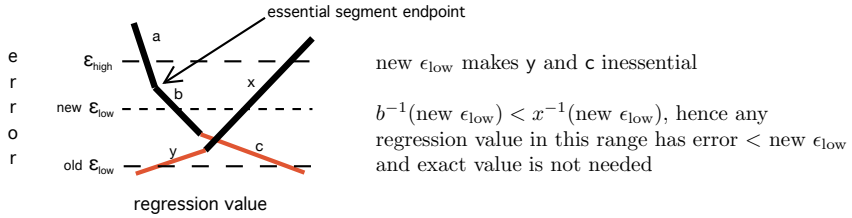

Figure 1: Downward and upward bounded envelopes

Theorem 1 *Given weighted data $(\mathbf{y}, \mathbf{w})$, sorted by the independent coordinate, and number of steps b, one can determine an optimal L_∞ b-step approximation in $\Theta(n + \log n \cdot b(1 + \log n/b))$ time and $\Theta(n)$ space.*

Given a set $(\mathbf{y}, \mathbf{w})$ of weighted values and $k \in [1\!:\!n]$, the *1-dimensional weighted k-center problem* is to find a set $S = \{s_1, \ldots, s_k\}$ of real numbers that minimizes $\max\{d(y_i, S) : i \in [1\!:\!n]\}$, where $d(\cdot, S)$ is the weighted distance to S, i.e., $d(y_i, S) = \min\{w_i \cdot |y_i - s_j| : j \in [1\!:\!k]\}$. Note that a set S is an optimal k-center iff it is the step values of an optimal L_∞ k-step approximation of the values in sorted order.

From now on we generally omit mention of "L_∞" and "optimal" since they are implied.

2.1 Interval Tree of Bounded Envelopes

For a weighted value (y, w), in the y-z plane the error of using $z \geq y$ as its regression value is given by the ray in the upper half-plane that starts at $(y, 0)$ with slope w. Given a set of weighted data $(\mathbf{y}, \mathbf{w})$, its *upward error envelope* is the topmost sequence of line segments corresponding to all such rays. For each z, it gives the maximum error of using z as the regression value for all points (y_i, w_i) where $z \geq y_i$. The *downward error envelope* uses rays in the upper half-plane starting at $(y_i, 0)$ with slope $-w_i$, representing the error of using a regression value $\leq y_i$. The intersection of the downward and upward error envelopes gives the regression value minimizing the error over the entire set, i.e., the weighted L_∞ mean, and its error.

To simplify exposition we assume that n is an integral power of 2. A binary *interval tree* has a root corresponding to the interval $[1\!:\!n]$, its two children correspond to $[1\!:\!n/2]$ and $[n/2+1\!:\!n]$, their children represent intervals of length $n/4$, etc. The leaves are $[1\!:\!1]$, $[2\!:\!2]$, $\ldots$ $[n\!:\!n]$. The intervals corresponding to nodes will be called *binary intervals*.

Some authors [2, 5, 9] used interval trees where each node contains the upward and downward error envelopes of the data in its interval, but most of the envelopes' segments are unnecessary. Let ϵ_{opt} be the (unknown) error of an optimal b-step approximation, and suppose bounds $\epsilon_{\text{low}} < \epsilon_{\text{opt}} \leq \epsilon_{\text{high}}$ are known. Then ϵ_{opt} can be determined using only the segments representing errors in $(\epsilon_{\text{low}}, \epsilon_{\text{high}})$. These will be called *essential segments*, and they form *bounded envelopes*. All others, the *inessential segments*, are discarded. Figure 1 shows how essential segments can become inessential when a better error bound is determined. Initially $\epsilon_{\text{low}} = 0$ and $\epsilon_{\text{high}} = \infty$, and the algorithm continually improves these bounds. In our interval tree each node contains its upward and downward bounded envelopes. Throughout, the essential segments are precisely the segments in the standard, unbounded, envelopes that represent errors in $(\epsilon_{\text{low}}, \epsilon_{\text{high}})$. Thus correctness depends on properties of the standard envelopes, though timing does not.

Bounded envelopes are stored as a doubly-linked list ordered by slope. Whenever a node is visited, by starting at both ends, inessential segments (i.e., those with no errors in $(\epsilon_{\text{low}}, \epsilon_{\text{high}})$) are discarded. The time is charged to the segments removed, not the search visiting the node. Only $\Theta(n)$ segments are ever created, hence the total time to remove inessential ones is $\Theta(n)$. Whenever the number of remaining segments is counted the count is only of the essential segments given the current values of ϵ_{low} and ϵ_{high}.

2.2 Feasibility Tests

Given an arbitrary interval let $U(S)$ denote the upward bounded envelope of the data in S, and $D(S)$ the downward bounded envelope. For $\epsilon \in (\epsilon_{\text{low}}, \epsilon_{\text{high}})$ the error of making S a single step is $<, =, > \epsilon$ iff $D(S)^{-1}(\epsilon) <, =, > U(S)^{-1}(\epsilon)$ (see Figure 1). Since $\epsilon \in (\epsilon_{\text{low}}, \epsilon_{\text{high}})$ this can be calculated as follows: let $S = \cup_{j=1}^{k} I_j$ for some $k \geq 1$, where each I_j is a binary interval. Then $U(S)^{-1}(\epsilon) = \min_{j=1}^{k} U(I_j)^{-1}(\epsilon)$ and $D(S)^{-1}(\epsilon) = \max_{j=1}^{k} D(I_j)^{-1}(\epsilon)$. For a binary interval I, to determine $U(I)^{-1}(\epsilon)$, and similarly $D(I)^{-1}(\epsilon)$, go through the segments of its bounded envelope until the segment r containing ϵ is found. This search alternates back and forth starting at the topmost and bottommost essential segments. This is only performed during a feasibility test, which will result in either the segments above r, or those below r, becoming inessential. Thus the time to find r is at most a constant plus a term linear in the number of segments that become inessential. Here too the linear term is charged to the inessential segments.

Any interval $[i : j]$ can be decomposed into $\leq 2\lfloor \lg(j-i+1) \rfloor + 1$ binary intervals where the sizes increase and then decrease, with perhaps two intervals of the same size in the middle. E.g., $[2 : 13] = [2\,{:}\,2] \cup [3\,{:}\,4] \cup [5\,{:}\,8] \cup [9\,{:}\,12] \cup [13\,{:}\,13]$. These can be visited in $O(\log n)$ time by a tree traversal starting at the leaf $[i : i]$, moving upward to the least common ancestor of i and j, and then downward to the leaf $[j\,{:}\,j]$. Suppose, given i and $\epsilon \in (\epsilon_{\text{low}}, \epsilon_{\text{high}})$, we want to find the largest j such that the error of making $[i\,{:}\,j]$ a single step is $\leq \epsilon$. We do this by a traversal to locate $j + 1$. By incrementally updating the $\min$ values to determine $U_S^{-1}(\epsilon)$, and $\max$ values used for $D_S^{-1}(\epsilon)$, when moving upward at node p one can determine if $j + 1$ is in p's subtree (and hence the traversal should start going downward) by using p's envelopes to decide if adding the entire subtree gives an error $> \epsilon$. When moving downward, $j + 1$ is in p's left subtree iff adding the left subtree gives error $> \epsilon$, otherwise it is in p's right subtree. Not counting the queries of children's envelopes, the nodes visited are the same as those in going from i to $j + 1$ when $j + 1$ is known.

Given b and $\epsilon \in (\epsilon_{\text{low}}, \epsilon_{\text{high}})$, a *feasibility test* determines if there is a b-step function with regression error $\leq \epsilon$. This can be accomplished by starting at 1 and determining the largest j_1 for which the error of making $[1\,{:}\,j_1]$ a single step is $\leq \epsilon$, then starting at j_1+1 and determining the largest j_2 for which the error of making $[j_1 + 1\,{:}\,j_2]$ a single step is $\leq \epsilon$, etc. If the b^{th} step is finished before n is reached then ϵ is infeasible, the test stops, and $\epsilon_{\text{low}} = \epsilon$. Otherwise, ϵ is feasible, the steps have been identified, and $\epsilon_{\text{high}} = \epsilon$.

To count the number of nodes visited, for each step the traversal visits nodes at a given height at most twice, once moving upward and once moving downward. Thus at any height at most $2b$ nodes are visited. The top $\lfloor \lg b \rfloor$ levels have a total of $\Theta(b)$ nodes. There are $\lceil \lg n \rceil - \lfloor \lg b \rfloor = \Theta(\log n/b)$ lower levels, so in total $\Theta(b(1 + \log n/b))$ nodes are visited. Each visit takes $\Theta(1)$ time, so this is also the time required.

2.3 Constructing the Tree

We reduce the time to construct the interval tree of bounded envelopes by continually shrinking $(\epsilon_{\text{low}}, \epsilon_{\text{high}}]$. At the end, $(\epsilon_{\text{low}}, \epsilon_{\text{high}}]$ is so small that each bounded envelope is a single segment. See Figure 2.

First a feasibility test with $\epsilon = 0$ is performed using only the base level. If it passes then the algorithm is done. Otherwise, set $\epsilon_{\text{low}} = 0$, $\epsilon_{\text{high}} = \infty$, and $R = \emptyset$. Throughout, R is an unordered multiset containing the errors of all segment endpoints (e.g., the error of the joint endpoint of a and b in Fig. 1) remaining in $(\epsilon_{\text{low}}, \epsilon_{\text{high}})$. At height 0 each interval is a singleton, with single rays in its upward and downward envelopes for a total of $2n$ rays. In general, at the end of height h, $|R| < n/2^h$ and the total number of essential segments in the envelopes at height h is $< 3n/2^h$. Moving upward, bounded envelopes from height h are merged to form those at height $h + 1$, creating $< 3n/2^h$ segments and $< 2n/2^h$ segment endpoints (the number of segment endpoints is the number of segments minus one per envelope). Add the errors of the segment endpoints to those already in R, resulting in $|R| < 3n/2^h$. Take the median error in R and do a

initialize envelopes of leaf nodes, $\epsilon_{\text{low}} = 0$, $\epsilon_{\text{high}} = \infty$, R $= \emptyset$
for h=1 to $\lg n$ {h is height}
    for every binary interval I at height h, make I's envelopes by merging children's envelopes,
        and add any segment endpoint errors in $(\epsilon_{\text{low}}, \epsilon_{\text{high}})$ to R
    repeat 3 times {reducing $|R|$ to $< n/2^h$ and total essential segments at height h $< 3n/2^h$}
        feasibility test using median of remaining essential segment endpoint errors in R
do 2 more feasibility tests, reducing R to $\emptyset$, i.e., all envelopes are single essential segments

Figure 2: Constructing the Tree of Bounded Envelopes, Feasibility Tests Continually Shrink $(\epsilon_{\text{low}}, \epsilon_{\text{high}}]$

feasibility test. Depending on the outcome, one of ϵ_{low} and ϵ_{high} is adjusted and 1/2 the entries in R can be eliminated. Doing this 3 times results in $|R| < n/2^{h+1}$. The number of essential segments is $\leq |R| + 1$ per envelope (since R included segment endpoints at level $h+1$), and hence is $< 3n/2^{h+1}$.

When the top is finished $|R| \leq 2$ and 2 feasibility tests are used to eliminate the remaining endpoint errors, i.e., at every node of the interval tree the upward and downward bounded envelopes have only one segment. Complete the tree construction by removing all inessential segments, taking $\Theta(n)$ time. Feasibility tests during tree construction have a slight change from standard traversals in that when height h is being constructed, when the test's traversals reach height h they go sideways, not upwards, from one node to the next since nodes at higher levels haven't yet been constructed. This increases the total number of nodes visited per test by at most $n/2^h$. Since only 3 tests are done per height (see Figure 2), this adds $\Theta(n)$ total time over all heights. Thus the total time to construct the tree is $\Theta(n + \log n \cdot b(1 + \log n/b))$.

2.4 Search for Minimal Feasible Error

The L_∞ error of a stepwise approximation is the maximum of the L_∞ errors of its steps, thus there is an interval $[i:j]$ such that the error of an optimal b-step approximation is the error of using the weighted L_∞ mean as the step value on $[i:j]$. Thus searching through such errors can determine the minimal feasible error. "Parametric search" was used in [6, 11] but this is only of theoretical interest since parametric search is completely impractical, involving very complex data structures and quite large constants.

Search in a sorted matrix provides a practical approach [7]. Let E be the $n \times n$ matrix where $E(i,j)$ is the error of using the L_∞ mean on $[i:j]$ if $i \leq j$, and is 0 if $i > j$. E is not actually created, but serves as a conceptual guide. Few of its entries are ever determined. Its rows are nondecreasing and the columns are nonincreasing, so for any submatrix its largest entry is in the upper right and the smallest is in the lower left.

The algorithm has stages $0 \ldots \lg n - 1$, where at the start of stage s there is a collection of disjoint square submatrices of size $n/2^s$. Stage 0 starts with all of E. At each stage, divide each of the matrices into quadrants, and let ϵ_1 be a median of the smallest value from each quadrant and ϵ_2 a median of their largest values. Values determined to be outside $(\epsilon_{\text{low}}, \epsilon_{\text{high}})$, as in Figure 1, are not calculated exactly and are set to $\epsilon_{\text{low}} - 1$ or $\epsilon_{\text{high}} + 1$, as appropriate. Feasibility tests for ϵ_1 and ϵ_2 are done, resulting in improvements to ϵ_{low} and/or ϵ_{high}. Quadrants with smallest value $\geq \epsilon_{\text{high}}$, or largest value $\leq \epsilon_{\text{low}}$, are eliminated. The remaining quadrants are the matrices that start the next stage. Note that if $\epsilon_{\text{opt}} < \epsilon_{\text{high}}$ then any quadrant with an entry of ϵ_{opt} is not eliminated, hence at the end of each stage, either $\epsilon_{\text{opt}} = \epsilon_{\text{high}}$ or one of the entries in the remaining submatrices is ϵ_{opt}.

After the last stage the remaining matrices are 1×1 and a standard binary search on these values is used to finish the determination of ϵ_{opt}. The search uses $\Theta(\log n)$ feasibility tests, and the less obvious fact, proven in [7], is that only $\Theta(n)$ entries of E are evaluated.

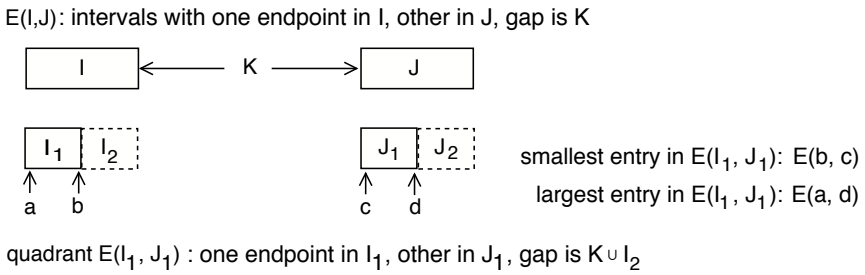

Figure 3: Evaluating smallest, largest entries in E(I_1, J_1) using gap from $E(I, J)$

2.5 Evaluating E During The Search

For intervals $I, J \subseteq [1:n]$ let $E(I, J)$ denote the submatrix $\{E(i, j) : i \in I, j \in J\}$, i.e., the submatrix corresponding to intervals starting at some $i \in I$ and ending at some $j \in J$. At the start of stage s of the search there is a collection of submatrices of the form $E(I, J)$ for binary intervals I, J of size $n/2^s$. Either $I = J$, or I is to the left of J and there is a (perhaps empty) gap between them with length an integral multiple of $n/2^s$ (K in Figure 3). During stage s, the quadrants of $E(I, J)$ are formed by cutting I and J in half into I_1, I_2 and J_1, J_2, respectively, creating quadrants $E(I_1, J_1)$, $E(I_1, J_2)$, $E(I_2, J_1)$, and $E(I_2, J_2)$. The smallest and largest value in each quadrant needs to be determined, and as Figure 3 shows, one can evaluate the smallest entry in $E(I_1, J_1)$ (i.e., $E(b, c)$), by using the envelopes from K and the binary intervals $[b:b]$, I_2, and $[c:c]$. The largest entry in the quadrant, $E(a, d)$, uses envelopes from K and the binary intervals I and J_1. Similar results hold for all of the other quadrants of $E(I, J)$. Exact values for entries outside $(\epsilon_{\text{low}}, \epsilon_{\text{high}}]$ are irrelevant and $\epsilon_{\text{low}} - 1$ or $\epsilon_{\text{high}} + 1$ is used, as appropriate.

The bounded envelopes for gap K are associated with $E(I, J)$, and if, say, $E(I_1, J_1)$ is kept for stage $s+1$ then the envelopes for $I_2 \cup K$ are associated with it. Just as for the tree construction, as search in a sorted matrix is proceeding the number of segments in the gap envelopes is reduced by interleaving feasibility tests based on the segment endpoint errors with tests for the basic search. At the end of stage s each gap envelope is copied at most 4 times, and each quadrant passed on to the next level adds ≤ 3 binary intervals which have envelopes that are only single segments. As shown in [7] there are $\leq 2^{s+3} - 1$ such quadrants, so by using 2 additional feasibility tests at each stage the total number of segments in the bounded envelopes at stage s is $O(2^s)$. Since the time to evaluate an entry of E is linear in the number of segments involved, the total time to evaluate entries of E over all stages of the algorithm is $\Theta(n)$, and the total time of the $\Theta(\log n)$ feasibility tests is $\Theta(\log n \cdot b(1 + \log n/b))$ plus the $\Theta(n)$ time to remove inessential segments.

This completes the proof of Theorem 1.

3 Final Comments

We have shown how to find an L_∞ b-step approximation of weighted data, presorted by its independent coordinate, in $\Theta(n + \log n \cdot b(1 + \log n/b))$ time. No previous algorithm [2, 3, 4, 5, 6, 8, 9, 11, 12, 13, 14] was $o(n \log n)$ whenever $b = o(n)$, nor $\Theta(n)$ whenever $b = O(n/\log n \log\log n)$. For sorted data the algorithm solves the 1-dimensional weighted k-center problem in the same time.

With a small change the algorithm also produces a "reduced isotonic" b-step function. f is an *isotonic* function iff $f(1) \leq f(2) \leq \ldots \leq f(n)$, and is an *optimal L_∞ b-step reduced isotonic regression of* $(\mathbf{y}, \mathbf{w})$ iff it minimizes the L_∞ error among all isotonic b-step functions. Isotonic regression is an important form of nonparametric regression that allows researchers to replace parametric assumptions with weaker shape

constraints [1, 15]. Some researchers were concerned that it can overfit the data and/or be too complicated [10, 16, 17] and resorted to reduced isotonic regression. However, they used approximations because previous exact algorithms were too slow. Merely changing the feasibility test to insure increasing steps finds b-step reduced isotonic regression in the same time bounds as b-step approximation.

Acknowledgement: Research partially supported by NSF grant CDI-1027192

References

[1] Barlow, RE, Bartholomew, DJ, Bremner, JM, and Brunk, HD, *Statistical Inference Under Order Restrictions: The Theory and Application of Isotonic Regression*, John Wiley, 1972.

[2] Chen, DZ and Wang, H, "Approximating points by a piecewise linear function", *Algorithmica* 66 (2013), pp. 682–713.

[3] Chen, DZ and Wang, H, "A note on searching line arrangements and applications", *Info. Proc. Let.* 113 (2013), pp. 518–521.

[4] Díaz-Báñnez, JM and Mesa, JA, "Fitting rectilinear polygonal curves to a set of points in the plane", *Eur. J. Operational Res.* 130 (2001), pp. 214–222.

[5] Fournier, H and Vigneron, A, "Fitting a step function to a point set", *Algor.* 60 (2011), pp. 95–109.

[6] Fournier, H and Vigneron, A, "A deterministic algorithm for fitting a step function to a weighted point-set", *Info. Proc. Let.* 113 (2013), pp. 51–54.

[7] Frederickson, G and Johnson, D, "Generalized selection and ranking: Sorted matrices", *SIAM J. Comp.* 13 (1984), pp. 14–30.

[8] Fülöp, J and Prill, M, "On the minimax approximation in the class of the univariate piecewise constant functions", *Oper. Res. Let.* 12 (1992), pp. 307–312.

[9] Guha, S and Shim, K, "A note on linear time algorithms for maximum error histograms", *IEEE Trans. Knowledge and Data Engin.* 19 (2007), pp. 993–997.

[10] Haiminen, N, Gionis, A, and Laasonen, K, "Algorithms for unimodal segmentation with applications to unimodality detection", *Knowl. Info. Sys.* 14 (2008), pp. 39–57.

[11] Hardwick, J and Stout, QF, "Optimal reduced isotonic reduction", *Proc. Interface 2012*.

[12] Karras, P, Sacharidis, D., and Mamoulis, N, "Exploiting duality in summarization with deterministic guarantees", *Proc. Int'l. Conf. Knowledge Discovery and Data Mining* (KDD) (2007), pp. 380–389.

[13] Liu, J-Y, "A randomized algorithm for weighted approximation of points by a step function", *COCOA* 1 (2010), pp. 300–308.

[14] Mayster, Y and Lopez, MA, "Approximating a set of points by a step function", *J. Vis. Commun. Image R.* 17 (2006), pp. 1178–1189.

[15] Robertson, T, Wright, FT, and Dykstra, RL, *Order Restricted Statistical Inference*, Wiley, 1988.

[16] Salanti, G and Ulm, K, "A nonparametric changepoint model for stratifying continuous variables under order restrictions and binary outcome", *Stat. Methods Med. Res.* 12 (2003), pp. 351–367.

[17] Schell, MJ and Singh, B, "The reduced monotonic regression method", *J. Amer. Stat. Assoc.* 92 (1997), pp. 128–135.